Evolved simulated agents exhibit size constancy abilities in solving an online size discrimination task

Massimiliano Schembri (massimiliano.schembri@uniroma1.it)
Department of Psychology, University of Rome "La Sapienza",
Via dei Marsi 78
00183, Rome, Italy

Marta Olivetti Belardinelli (marta.olivetti@gmail.com)
ECONA, Interuniversity Centre for the Research on Cognitive Processing in Natural and Artificial Systems,
Rome, Italy

Abstract

We describe some Artificial Life simulations in which a situated model agent controlled by a feed-forward neural network has to solve a simple categorization task involving size constancy abilities in an online fashion. The results show that even a simple neural controller without internal recurrent dynamics is capable of solving a non-trivial size categorization task by exploiting the dynamical interaction of the agent with its environment. Even if at an early stage, this work suggests two possible implications for the study of size constancy and perceptual constancy in general. First, approaching the problem from a functional point of view may open new perspectives on the possible underlying mechanisms. Second, the adoption of an embodied and situated approach may help to explain why perceptual constancy is so efficient in biological cognitive systems.

Keywords: perceptual constancy, size constancy, active perception, dynamical categorization

Introduction

Perceptual constancy can be either defined as a perceptual mechanism or as a behavior. In the first case we define it as the ability to perceive the stable properties of the surrounding environment despite the continuous change of the raw information reaching our sense organs. The second type of definition can be expressed saying that perceptual constancy allows us to behave in accordance with the stable properties of the surrounding environment. One definition puts emphasis on the mechanism (perception) and the other one on the behavior, but, as claimed by Ittelson (1951), "Any complete theory of perceptual constancy must encompass all its aspects" and therefore should consider both the mechanisms and the behaviors. Instead, most of the recent theories and computational models of perceptual constancy focus on the presumed underlying mechanisms, but tell us very little about how they translate into functional behaviors. Some examples of constancy mechanisms proposed in the literature are mental rotation (Jolicoeur & Humphrey, 1998), perceptual compensation (Bridgeman, 2010), 3D reconstruction (Edelman & Weinshall, 1998) and hierarchical feature extraction (Foldiak, 1998).

The behavioral aspects of perceptual constancy, instead, are hugely neglected except for certain animal research studies where constancy also reveals its great ecological relevance. Size-dependent food selection, evaluation of predator size and distance, foraging in different daylight conditions are some example of behaviors that have been studied showing some constancy abilities (but also failures) even in lower vertebrates. Size constancy has been studied in frogs and toads for example by Ingle (1968), Ingle and Cook (1977), Lettvin, Maturana, McCulloch and Pitts (1959). Shape invariance has been studied in fishes and amphibians (e.g. Ingle, 1963; Ingle, 1971; Ewert, 1984). A lot of research on color constancy has been conducted with experiments on bees, amphibians, fishes, cats and monkeys (Neumeyer, 1998).

The approach proposed here is based on the methodology of Artificial Life (Langton, 1998) and Evolutionary Robotics (Nolfi, 1998; Nolfi & Floreano, 2000) and is an attempt to build a minimal but complete model of size constancy capable of simulating a functional behavior and able to explain some aspects of the cognitive mechanisms of size constancy. The general idea at the base of this approach is that perceptual constancy cannot be properly understood studying a cognitive system in isolation and detached from its natural context. It seems to be a research area in which an embodied and situated approach is essential. This idea is not completely new, since there have been some experiments on size constancy with an Evolutionary Robotic approach that started to envision the problem with an embedded and situated approach (Scheier, Pfeifer, Kunyioshi, 1998; Nolfi & Marocco, 2000). More recently Williams & Beer (2010) proposed some simulations in which a simulated model agent is evolved to discriminate between small and big circles. The work described here shares the same approach but uses different kind of sensors and actuators and an online task (not based on single trials). More in general, our goal is to develop an embodied and situated framework for studying different aspects of size constancy in a systematic way and from a functional perspective, and the results described in this work seem to support this endeavor.

Methods

The experimental setup proposed here is a computer simulation that represents a simplified model of a brain-body-environment system with the following characteristics:

1) A simulated agent with a sensory-motor system acts in a virtual environment
2) Sensory input and its variations are coherent with the environment structure and its laws
3) Variation of the input is partly determined by the motor system
4) The neural controller of the agent evolves through a Genetic Algorithm, with no prior hypothesis about its functioning
5) The fitness function used to evolve the neural controller is based on the agent performance in a task that requires some degree of perceptual constancy

The main goal of this experimental setup is to provide an embodied and situated context in which a simulated agent can evolve a size constancy behavior.

Simulation Environment

The simulation environment is described in figure 1. A simulated agent moves in a 2D square arena with sides of length 60 populated with circles randomly placed in a grid of 5x5 cells positions (figure 1 top part). The diameters of the circles can be small (0.5) or big (1.0). There are 10 small and 10 big circles for a total amount of 20 objects.

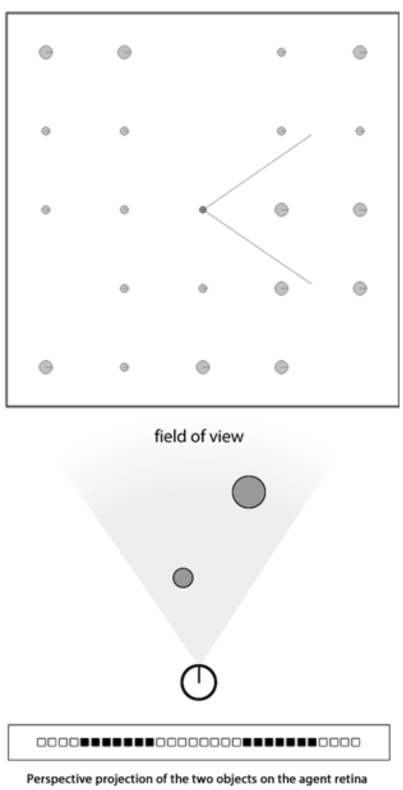

Figure 1: Experimental Setup

The agent (see figure 1 bottom part), represented by a small circle of size 0.5, is provided with a linear array of visual receptors by which it is able to "see" objects in front of it with a field of view of 60°. The activation of the receptors is calculated with a perspective projection of the objects in the field of view of the agent so that a near small circle and a distant one can have the same retinal projection (as depicted in figure 1). Distance cues are provided through a sort of "fog effect" (not shown in figure 1) that makes the circles appear lighter and lighter as the viewing distance increases. The fog effect and the grid configuration of objects make the agent input clean and avoid cluttered input patterns. The fog effect in particular avoids that too many objects are visualized at the same time on the retina. This would require the agent to develop some kind of attentional mechanism that would deserve a dedicated work.

The controller of the agent is a three layer feed-forward neural network. The input layer is the above mentioned linear retina with 30 receptors whose activations range from 0.0 to 1.0. The hidden layer has 10 units and the motor layer consists of 4 output units. Both hidden and motor layer neurons use a sigmoid activation function. Each layer is fully connected with the next one. So there are 300 input-hidden weights (30x10) and 40 hidden-output weights (10x4). Figure 2 shows the structure of the neural network.

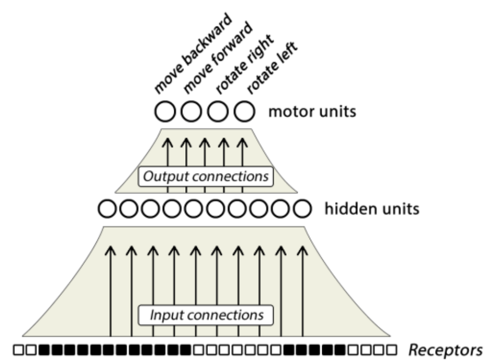

Figure 2: Neural Controller of the robot

The agent can move forward or backward with a certain velocity and can rotate around its center to change direction. The four motor units control the movement of the agents with two couples of opposing real units. The linear movement of the robot is determined by the results of two opposing units that determines the forward and backward linear velocities. The agent moves forward if the output of the forward velocity unit is higher than the backward one, and vice versa. The agent direction is determined by two opposing units controlling the right and left angular velocity.

Task

The goal of the agent is to hit as much small circles as possible and to avoid the big ones during its lifetime that lasts 10,000 simulation steps. Circles that are hit by the agent are removed from the environment.

Evaluating circle size is not a trivial task because during environment exploration the sensory input varies continuously and produces ambiguous configurations. The same retinal projection, for example, can be that of a near small circle or the one of a big distant one. So the retinal subtense in itself is not correlated with object size. The same occurs for the object "lightness" that varies with distance. The organism faces a size constancy problem.

The task is similar to the one proposed by Scheier, Pfeifer, Kunyioshi (1998) and Nolfi and Marocco (2000) and more recently by Williams and Beer (2010), but the motor system proposed here is different allowing for fast forward and backward linear movements. Moreover with respect to the work of Nolfi and Marocco (2000) and Williams and Beer (2010) the task is not based on single separate trials but requires an online behavior in which the single discriminations occur seamlessly during the entire life of the robot without resetting the experimental setup after each robot response.

Genetic Algorithm

A genetic algorithm is used to evolve the weights of the neural network to solve the simple size discrimination task described above. As mentioned before, the goal of the agent is to hit as much small objects as possible and to avoid the big ones. Objects that are hit by the agent are removed from the environment. The fitness function is calculated with the following formula:

$$\mathbf{F = C_s - C_b}$$

where C_s and C_b are the number of small and big circles hit at the end of the agent life. Since there are a total of 10 small circles and 10 big ones, the highest fitness score is 10 and the lowest is -10. The evolutionary experiment consists in evolving the weights of the neural controllers in a populations of 100 agents for 100 generations with a selection criterion based on the fitness function described before.

The weights of the neural networks in the first generation are initialized in the range (-1/sqrt(d), +1/sqrt(d)) where d is the number of input to each neuron. When all the individuals of one generation have been tested they are sorted based on their fitness scores and the 20% of the best individuals are selected to produce the next generation of agents. The genetic operator consists of a mutation mechanism that changes 10% of the weights of the neural network adding a random number between -0.5 and +0.5. The genetic algorithm uses elitist selection allowing the best individual of one generation to carry over to the next generation with unaltered connection weights.

Results

The genetic algorithm described above was used to run 10 seeds of the same simulation some of which gave interesting results. Figure 3 shows the best and mean fitness along each generation of the best simulation obtained. The best fitness of the best individual in the last generation is 9, which is nearly the maximum score possible for the fitness. This result indicates that the evolution process produced some kind of behavior capable of avoiding big circles and approaching and hitting the small ones.

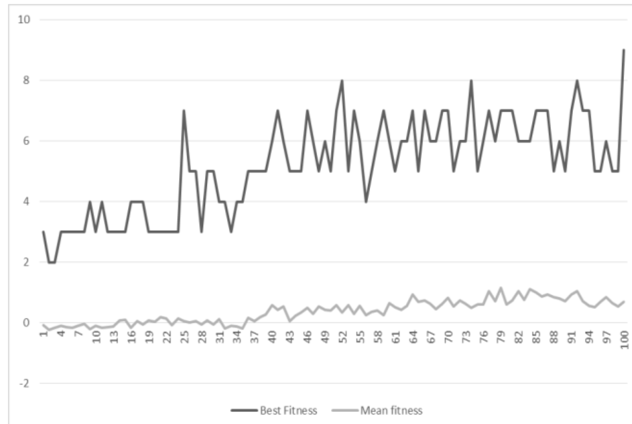

Figure 3: Graph of the best and mean fitness for each of the 100 generations of the best run

Considering that, as explained before, the size of the circle cannot be evaluated relying on a single retinal projection at a given moment, or the intensity of retinal receptors, we can expect the evolved neural system to develop a form of size constancy behavior based on the dynamical interaction of the agents with its environment to exploit the information contained in the optic flow as theorized by Gibson (1950/1966) and demonstrated in some classical research studies(Lee, 1980; Franceschini et Al. 1992).

It could sound strange that a simple neural controller like the one used in this experiment is capable of such complex behaviors. Actually a feed-forward neural network is a simple type of controller with no internal states, where information flows in only one direction, with each input producing always the same output. In this respect it is comparable with a simple associative mechanism. What makes this experiment interesting is that the neural network is inside an embodied and situated agent whose sensor and motor systems allow it to interact with its environment (Figure 4). Each input, at a given time, produces an output that is used to move the agent. The agent movement, in turn, changes the next input, which produces a new output and so on. This mechanism gives rise to interesting "organism-environment" dynamics that allow the agent exhibit a functional size discrimination behavior.

Some preliminary behavior analysis have been performed on the best organism of the last simulation and gave some interesting results. First of all, to accomplish their task, most of the successful organism develop a sort of exploratory behavior consisting in turning around their centers and moving slowly until some object fall in their receptive field. Once an object shows up in the receptive field the agent gets close to the object and then start to oscillate back and forth for a few times.

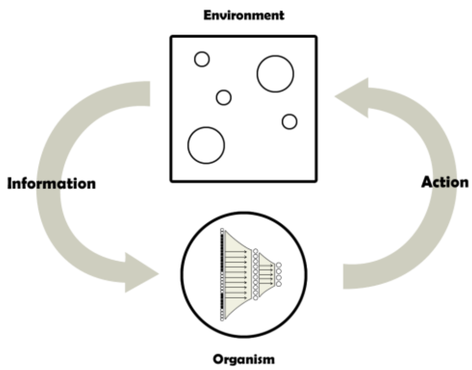

Figure 4: Organism-Environment relationships

At this point the behavior is different depending on whether the object is a small circle or a big one. In the case of small circles the agent goes forward and hit it (see figure 5).

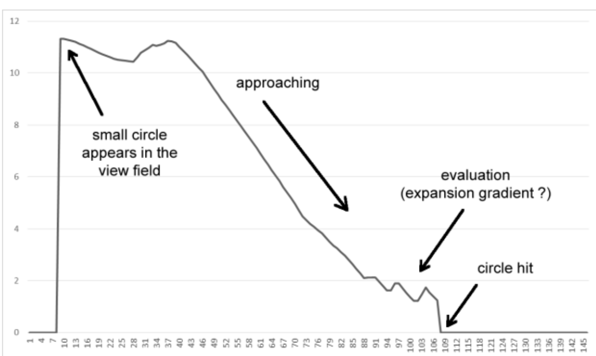

Figure 5: Graph of the interaction with small circles (y-axis = distance from the object, x-axis = simulation steps)

In the case of big circles the oscillating behavior ends with the agent getting away from the object (see figure 6) towards a location favorable for the complete exploration of the environment. The oscillating behavior could be interpreted as a discrimination phase and always takes place at approximately the same distance (about 2.0) from the target object.

At a first glance it could be thought that a simple "hand – made" linear function using the number and intensity of receptors should be enough to discriminate between the large and small circles. But looking at the results and considering the dynamical context in which the agent lives it is clear that a far more complex behavior is required to solve the task. Indeed, the behavior obtained with the genetic algorithm is quite articulated and comprises at least five sub components: explore, approach, discriminate, hit object, avoid object. Moreover, each of this sub components of the behavior has a time course and therefore is more complex than a one shot discrete response. This should be enough to convey the idea that designing by hand a system capable of acting in a dynamical environment is not a trivial task.

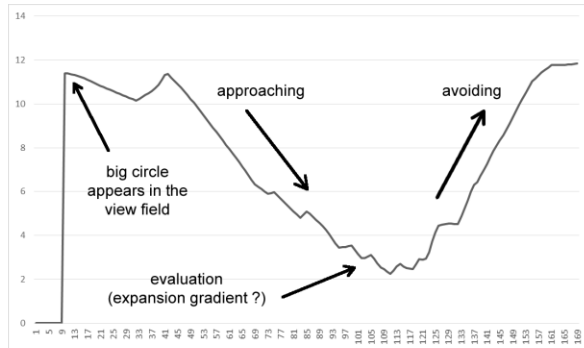

Figure 6: Graph of the interaction with big circles (y-axis = distance from the object, x-axis = simulation steps)

Further analysis are required to better understand what happens inside the neural controller and to explain the agent behavior in more detial. The most tempting hypothesis, at the moment, is that the agent performs some kind of expansion gradient assessment during the discrimination phase as suggested by the fact that the oscillating behavior occurs more or less at the same distance from the object, and rather close to it. Indeed, the expansion gradient of two objects must be evaluated at the same distance, and the nearer an object is to the observer the wider and more informative its expansion gradient is. Some "laboratory" manipulation are needed to clarify this and many other aspects. For example we don't know how robust this behavior is in different environmental conditions, what happens if the agent starting position is changed or if the objects are not placed in a grid pattern.

Conclusions and future work

We described an Artificial Life simulation in which a simulated agent controlled by an evolved neural network shows some size constancy abilities in solving a simple categorization task. Even if a more detailed analysis is required, the preliminary results described here seem to confirm that a simple feed-forward sensory-motor system can solve a rather complex size constancy problem exploiting its dynamical interaction with the surrounding environment. These results strongly support an embodied-situated approach to perceptual constancy, and also suggest that the ability of a cognitive system can be better understood in a framework that fully considers the importance of the brain-body-environment dynamics.

In the future work we are planning to explore different experimental conditions varying the size constancy task, the agent sensory-motor apparatus and its neural controller.

References

Bridgeman, B. (2010). Space Constancy. The rise and fall of perceptual compensation. In: Space and time in perception and action. Nihawan, Romi (Ed.); Khurana, Beena (Ed.); New York, NY, US: Cambridge University Press, 2010. pp. 94-108.

Cliff, D. (1991). Computational Neuroethology: A Provisional Manifesto in J.-A. Meyer and S. W. Wilson, editors: From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour (SAB90). MIT Press Bradford Books, 1991, pp.29-39.

Cliff, D. (2003). Neuroethology, Computational". In: M. A. Arbib, editor: The Handbook of Brain Theory and Neural Networks. Second Edition. MIT Press Bradford Books. pp. 737-741

Edelman, S., & Weinshall, D. (1998). Computational approaches to shape constancy in Perceptual Constancy:Why things look as they do, eds. V Walsh & J J Kulikowski, Cambridge, U.K.: Cambridge Univ. Press , pp. 144-172, 1998

Ewert, J-P. (1984). Tectal mechanisms that underlie prey-catching and avoidance behaviors in toads. In H. Vanegas (Ed.), Comparative neurology of the optic tectum (pp. 247-416). New York:Plenum.

Franceschini, N., Pichon, J-M., & Blanes, C. (1992). From insect vision to robot vision. Philos Trans R Soc Lond B 337:283–294

Foldiak, P. (1998). Learning constancies for object perception, in Perceptual Constancy: Why things look as they do, eds. V Walsh & J J Kulikowski, Cambridge, U.K.: Cambridge Univ. Press, pp. 144-172, 1998

Gibson, J. J. (1950). The perception of the visual world. Boston:Houghton Mifflin.

Gibson, J. J. (1966). The Senses Considered as Perceptual Systems. Boston:Houghton Mifflin.

Ingle, D. (1963). Limits of visual transfer in goldfish. Unpublished Ph.D thesis, department of Psychology, University of Chicago. (as cited in Ingle, D. (1998))

Ingle, D. (1968). Visual releasers of prey-catching behavior in frogs and toads. Brain Behavior and evolution, 500-518

Ingle, D. (1971). The experimental analysis of visual behavior. In W.S. Hoar & D. J. Randall (Eds.) Fish physiology (Vol. 5, pp. 9-71). New York: Academic Press.

Ingle, D., & Cook, J. (1977). The effect of viewing distance upon size-preference of frogs for prey. Vision Research, 17, 1009-1014

Ingle, D. (1998). Perceptual Constancies in Lower Vertebrates. In V. Walsh, & J. J. Kulikowski (Eds.) Perceptual constancy: why things look as they do, (pp. 173–191). Cambridge University Press.

Ittelson, W. H. (1951). The constancies in perceptual theory. Psychological Review, 1951, 58, 285 - 294

Jolicoeur, P., & Humphrey, G. K. (1998). Perception of rotated two-dimensional and three-dimensional objects and visual shapes. In: V. Walsh, J. Kulikowski (eds) Perceptual constancy: why things look as they do. Cambridge University Press, Cambridge, pp.323-351

Kato T., & Floreano, D. (2001). An evolutionary active-vision system. In: Proceedings of the congress on evolutionary computation (CEC01), Seoul, Korea, October 2001. IEEE Press, Piscataway, NJ

Langton, C. G. (1998). Artificial life: an overview. MIT Press

Lee, D. N. (1980). The optic flow field: The foundation of vision. Philosophical Transaction of the Royal Society of London. Series B, 290, 169-179.

Lettvin, J. Y., Maturana, H. R., McCulloch, W. S., & Pitts, W. H. (1959). What the frog's eye tells the frog's brain. Proceedings of the IRE, 47, 1949-1951.

Mondada, F. & Floreano, D. (1995). Evolution of neural control structures: some experiments on mobile robots. Robotics and Autonomous Systems, 16:183-195.

Neumeyer, C. (1998). Comparative aspects of color constancy. In: V. Walsh, J. Kulikowski (eds) Perceptual constancy: why things look as they do. Cambridge University Press, Cambridge, pp.323-351

Nolfi, S. (1998). Evolutionary Robotics: Exploiting the full power of self-organization. In N. Sharkey (Ed.), Proceedings of Self-Learning Robots II: Bio-robotics. London: IEE Press, 1-7

Nolfi, S. & Marocco, D. (2000). Evolving visually-guided robots able to discriminate between different landmarks. In: J-A Meyer, A. Berthoz, D. Floreano, H.L. Roitblat, and S.W. Wilson (eds.) From Animals to Animats 6. Proceedings of the VI International Conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press. pp. 413-419

Nolfi, S. & Floreano, D. (2000). Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. Cambridge, MA: MIT Press/Bradford Books.

Scheier, C., Pfeifer, R., & Kunyioshi, Y. (1998). Embedded neural networks: exploiting constraints. Neural Networks, 1998, 11, 1551-1569

Williams, P. & Beer, R. (2010) Information Dynamics of Evolved Agents. In S. Doncieux, B. Girard, A. Guillot, J. Hallam, J.-A. Meyer and J-B. Mouret (Eds.), From Animals to Animats 11: Proceedings of the 11th International Conference on Simulation of Adaptive Behavior (pp. 38-49). Springer